arXiv:1402.7344v1 [cs.LG] 28 Feb 2014

An Incidence Geometry approach to Dictionary Learning*

Meera Sitharam
University of Florida

Mohamad Tarifi
Google

Menghan Wang
University of Florida

March 3, 2014

Abstract

We study Dictionary Learning (aka sparse coding). By geometrically interpreting an exact formulation of Dictionary Learning, we identify related problems and draw formal relationships among them. Dictionary Learning is viewed as the minimum generating set of a subspace arrangement. This formulation leads to a new family of dictionary learning algorithms. When the data are sufficiently general and the dictionary size is sufficiently large, we completely characterize the combinatorics of the associated subspace arrangements (i.e, their underlying hypergraphs). This characterization is obtained using algebraic and combinatorial geometry. Specifically, a combinatorial rigidity-type theorem is proven that characterizes the hypergraphs of subspace arrangements that generically yield (a) at least one dictionary (b) a locally unique dictionary (i.e., at most a finite number of isolated dictionaries) of the specified size. We are unaware of prior application of combinatorial rigidity techniques in the setting of Dictionary Learning, or even in machine learning. We list directions for further research that this approach opens up.

1 Introduction and Contributions: Geometric Dictionary Learning

Dictionary Learning is the problem of obtaining a sparse representation of data points, by learning dictionary vectors upon which the data points can be written as sparse linear combinations.

Definition 1. *Given an input data set $X = [x_1 \dots x_m]$, where $x_i \in \mathbb{R}^d$, and a required* sparsity s, Dictionary Learning *finds a* dictionary $D = [v_1 \dots v_n]$,*and* $\Theta = [\theta_1 \dots \theta_m]$, *where $v_j \in \mathbb{R}^d$ and $\theta_i \in \mathbb{R}^n$, such that $x_i = D\theta_i$ and $\|\theta_i\|_0 \leq s$. We then say that the dictionary D* s-represents X.

*This research was supported in part by the research grant NSF CCF-1117695 and a research gift from SolidWorks.

Specifically, we understand the problem from an intrinsically geometric point of view. This leads to a new class of algorithms and learning complexity theorems. Our *first contribution* is as follows:

- We cast Dictionary Learning in a noiseless, geometric setting, identify related problems and draw formal relationships among them. This leads to a view of dictionary learning as the minimum generating set of a subspace arrangement, and a new class of algorithms which applies subspace clustering techniques and intersection algorithms to learn the dictionary.

Notice that each $x \in X$ lies in an s-dimensional subspace spanned by vectors $v \in D$ that form the *support* of x, denoted $supp_D(x)$. The resulting *s-subspace arrangement* $S_{X,D} = \langle supp_D(x) : x \in X \rangle$ has an underlying (multi)hypergraph $H(S_{X,D}) = (\mathfrak{I}(D), \mathfrak{I}(S_{X,D}))$, where $\mathfrak{I}(D)$ denotes the index set of the dictionary D and the (multi)hyperedge set $\mathfrak{I}(S_{X,D})$ consists of the index sets corresponding to the sets $supp_D(x)$. The word "multi" appears because if $supp_D(x_1) = supp_D(x_2)$ for data points $x_1, x_2 \in X$ with $x_1 \neq x_2$, then that support set of dictionary vectors (resp. their indices) is multiply represented in $S_{X,D}$ (resp $\mathfrak{I}(S_{X,D})$). Thus we denote the sizes of these multisets as $\|S_{X,D}\|$ (resp. $\|\mathfrak{I}(S_{X,D})\|$), while the sizes of these sets after removing copies are denoted as usual as $|S_{X,D}|$ (resp. $|\mathfrak{I}(S_{X,D})|$).

We are interested in minimizing $|D|$ for general X. However, as a *corollary of the main result* (second contribution of this paper), we obtain the following.

- If the data points in X are highly general, for example, picked uniformly at random from the sphere S^{d-1}, then when s is fixed, $|D| = \Omega(|X|)$ with probability 1 (see Corollary 14).

Very little is known for simple restrictions on X. For example the following question is open.

Question 1. *Given a general position assumption on X, are smaller dictionaries possible than indicated by the above lower bound? Conversely, what is the best lower bound on $|D|$ under such an assumption?*

Question 1 gives rise to the following pure combinatorics open question.

Question 2. *Given weights $w(S) \in \mathbb{N}$ assigned to size-s subsets S of $[n]$. For $T \subseteq [n]$ with $|T| \neq s$,*

$$w(T) = \begin{cases} 0 & |T| < s \\ \displaystyle\sum_{S \subset T, |S| = s} w(S) & |T| > s \end{cases}$$

Assume additionally the following constraint holds: for all subsets T of $[n]$ with $s \leq |T| \leq d$, $w(T) \leq |T| - 1$. Can one give a nontrivial upper bound on $w([n])$?

Hence, more assumptions on X and D are required. Consider modeling the data X from a generative model. There are a few choices for generative models that produce data readily modeled by a dictionary. In its most general form we are asked to determine an unknown dictionary D and a set of unknown coefficients $\Theta = \{\theta_1 \dots \theta_m\}$,

given a set of sample points $X = \{x_1 \dots x_m\}$ such that $x_i = D\theta_i$ where $\|\theta_i\|_0 \leq s$. A further complication arises in the form of noise $x_i = D\theta_i + \epsilon_i$ where the l_2 norm of ϵ_i is bounded.

Alternatively, after learning D, if we are given a new data point x that is known to have a s-sparse representation over D (i.e. x lies on the s-subspace arrangement $S_{X,D}$), we would like to minimize the complexity of the *vector selection problem* of finding $supp_D(x)$, which depends not only on $|D|$, but also on the characteristics of D.

We say that a set V of vectors *s-spans* a point or a subspace if the point or all points in the subspace can be written as a linear combination of at most s elements of V. A common property often imposed on dictionaries is s-regularity: a dictionary D is *s-regular* if for all θ such that $\|\theta\|_0 \leq s$, it holds that $D\theta \neq 0$.

For an s-regular dictionary, the general vector selection problem is ill defined. For instance, D can be overcomplete, leading to multiple solutions for θ_i. Overcoming this by framing the problem as a minimization problem makes the problem exceedingly difficult. Indeed under generic assumptions, even determining the minimum l_0 norm of θ_i when D and x_i are *known* is NP-hard.

Under this condition, we can make the vector selection problem well defined by enforcing $2s$-regularity property on D. A dictionary D is *s-independent* if for all θ_1, θ_2, such that $\|\theta_1\|_0 \leq s$ and $\|\theta_2\|_0 \leq s$, it holds that $D\theta_1 = D\theta_2$ if and only if $\theta_1 = \theta_2$. The property s-independence is a minimal requirement for unique invertibility. Notice that the definition given for s-independence is indeed equivalent to $2s$-regularity.

We can further strengthen the constraints on D by assuming that D is a *frame*, that is, for all θ such that $\|\theta\|_0 \leq s$, there exists a δ_s such that $(1-\delta_s) \leq \frac{\|D\theta\|_2^2}{\|\theta\|_2^2} \leq (1+\delta_s)$. This ensures that basic tasks, such as vector selection, are tractable and noise tolerant.

Note: Unless mentioned otherwise, for the remainder of the manuscript, we set $\epsilon = 0$; i.e, we study *Exact Dictionary Learning.* We anticipate that our Exact Dictionary Learning results straightforwardly generalize to the approximate or noisy version of Dictionary Learning, where $\epsilon > 0$.

We impose a systematic series of increasingly stringent constraints on D and X that lead to a whole set of independently interesting geometric problems.

Problem 1 (Geometric Dictionary Learning)**.** *Let X be a given set of data points in $\mathbb{R}^d$, which is known to be s-spanned by a (optionally constrained to be frame) dictionary D with size $|D| \leq n$, using a subspace arrangement $S_{X,D}$ with $|S_{X,D}| \leq m$ i.e. $x_i = D\theta_i$, $\|\theta_i\|_0 \leq s$.* Geometric Dictionary Learning *is the problem of finding any (optionally frame) dictionary $\acute{D}$ satisfying the properties of D, i.e. $|\acute{D}| \leq n$, and exists $\acute{\theta}_i$ such that $x_i = \acute{D}\acute{\theta}_i$ for all $x_i \in X$.*

As a *corollary to our main result* (the third contribution of this paper), we obtain the following.

- An algorithm for Geometric Dictionary Learning, for sufficiently general data X, i.e., requiring sufficiently large dictionary size n (see Corollary 13).

Note that there could be many dictionaries D and for each D, many possible subspace arrangements $S_{X,D}$ that are solutions to the Geometric Dictionary Learning problem above.

As the *main result* (the fourth contribution of this paper) we use combinatorial rigidity techniques to obtain the following. To the best of our knowledge, this paper pioneers the use of combinatorial rigidity for Dictionary Learning.

- A complete characterization of the subspace arrangement hypergraphs $H(S_{X,D})$ that generically yield (a) at least one solution dictionary D (b) a locally unique solution dictionary D (i.e., at most a finite number of isolated solution dictionaries) of the specified size (see Theorem 8).

This leads to the following open question for future work.

Question 3. *What is the minimum size of a data set X such that the Geometric Dictionary Learning Problem for X has a unique solution dictionary D and associated subspace arrangement $S_{X,D}$? What are the geometric characteristics of such an X?*

One approach for the Geometric Dictionary Learning problem is decomposing it into subproblems. We define these problems below and then show their relationship to the original question. We additionally show how to combine algorithms for the subproblems into an algorithm for Geometric Dictionary Learning. These subproblems are also revealed to be of independent interest.

First, we introduce the notion of *support-equivalence* of data points in X with respect to a dictionary D. For a given subspace t in the subspace arrangement $S_{X,D}$ (respectively hyperedge h in the hypergraph's edge-set $\mathfrak{I}(S_{X,D})$), let $X_t = X_h \subseteq X$ be the equivalence class of data points x such that $span(supp_D(x)) = t$. We call the data points x in any given X_h as *support-equivalent*.

Problem 2 (Geometric Dictionary Learning for Partitioned Data)**.** *Given data X partitioned into $X_i \subseteq X$, find a dictionary D and s-subspace arrangement $S_{X,D}$ satisfying $|D| \leq n$, and $|S_{X,D}| \leq m$, such that X_i represent the support-equivalence classes of X with respect to D.*

A direct generalization of the above problem is the following. We say that a set of data points X *lies on* a set S of s-dimensional subspaces if for all $x_i \in X$, there exists $S_i \in S$ such that $x_i \in S_i$.

Problem 3 (Subspace Arrangement Learning)**.** *Let X be a given set of data points that are known to lie on a set S of s-dimensional subspaces of $\mathbb{R}^d$, where $|S|$ is at most s. (Optionally assume that the subspaces in S have bases such that their union is a frame).* Subspace arrangement learning *finds any subspace arrangement $\acute{S}$ of s-dimensional subspaces of $\mathbb{R}^d$ satisfying these conditions, i.e. $|\acute{S}| \leq s$, X lies on $\acute{S}$, (and optionally the union of the bases of $\acute{S}_i \in \acute{S}$ is a frame).*

The next of our key problems can be viewed as obtaining a minimally sized dictionary from a subspace arrangement.

Problem 4 (Smallest Spanning Set for Subspace Arrangement)**.** *Let S be a given set of s-dimensional subspaces of $\mathbb{R}^d$ specified by giving their bases (optionally their union is a frame). Assume their intersections are known to be s-spanned by a set I of vectors with $|I|$ at most n. Find any set of vectors $\acute{I}$ that satisfies these conditions.*

In general, the smallest spanning set is not necessarily unique even for s-regular dictionaries.

1.1 Problem Relationships and High Level Algorithms

Using solutions to the problems in the previous section, we give a two-step procedure that solves the s-regular Geometric Dictionary Learning problem, thereby clarifying how the problems are related.

- Learn a Subspace Arrangement S for X (instance of Problem 4).
- Recover D by finding the smallest Spanning Set of S (instance of Problem 3).

Note that it is not true that the decomposition strategy should always be applied for the same sparsity s, the constant in the generative model. The decomposition starts out with the minimum given value of s and is reapplied with iteratively higher s if a solution has not be obtained.

Furthermore, we observe that, for s-independent sets, the smallest spanning set can be obtained via intersection of the subspace arrangements. Under the condition that the subspace arrangement comes from an s-independent dictionary, the smallest spanning set is the union of: (a) the smallest spanning set I of the pairwise intersection of all the subspaces in S; (b) any points outside the pairwise intersections that, together with I, completely s-span the subspaces in S.

1.2 Cluster and Intersect Algorithm

Consider P_Θ, a distribution from which the coefficient vector θ_x for a data point x is generated. The *support* of θ_x is simply $supp_D(x)$. In this section P_Θ is as follows: (a) a set of k supports are picked uniformly from the set of 2^n possibilities; (b) the values of the non-zero entries of θ_i are picked uniformly from $\mathbb{R}^s$. This allows us to quantify the approach in terms of the number of subspaces k, which could vary between the various settings encountering an instance of the Dictionary Learning problem. For instance D is often used to separate causes in an environment, and naturally not all possible combinations of causes are realized.

There are several known algorithms for learning subspace arrangements. For a survey the reader is referred to [21]. Random Sample Consensus (RANSAC) [21] is an approach to learning subspace arrangements that isolates, one subspace at a time, via random sampling. When dealing with an arrangement of k s-dimensional subspaces, for instance, the method samples $s+1$ points which is the minimum number of points required to fit an s-dimensional subspace. The procedure then finds and discards inliers by computing the residual to each data point relative to the subspace and selecting the points whose residual is below a certain threshold. The process is iterated until we have k subspaces or all points are fitted. RANSAC is robust to models corrupted with outliers. The following algorithm illustrates random and deterministic RANSAC used for the present problem.

In our case, given a subspace arrangement S arising from an s-independent dictionary, the smallest spanning set can be written recursively in terms of the union of: (a) the spanning set of the arrangement, obtained by taking union of pairwise intersection of all the subspaces in S, together with points (b): outside the pairwise intersections that are would be necessary and sufficient to completely s-span the subspaces in S.

This directly leads to a recursive algorithm for the smallest spanning set problem. The dictionary is now obtained from picking m atoms from the intersection of the subspaces and the remaining spanning sets.

1.3 Complexity of Geometric Dictionary Learning

For Dictionary Learning, even the general vector selection problem of recovering Θ given X where D is *known* has been shown to be NP-hard by reduction to the Exact Cover by 3-set problem [14]. One is then tempted to conclude that Geometric Dictionary Learning is also NP-hard. However, this cannot be directly deduced in general. The error in this reasoning is that, even though adding a witness D turns the problem into an NP-hard problem, it is possible that the Geometric Dictionary Learning solves to produce a different dictionary $\acute{D}$.

1.4 Review: Traditional and Statistical Approaches to Dictionary Learning

An optimization version of Dictionary Learning can be written as:

$$\min_{D \in \mathbb{R}^{d,n}} \max_{x_i} \min \|\theta_i\|_0 : x_i = D\theta_i.$$

In practice, the Dictionary Learning problem is often relaxed to the Lagrangian $\min \sum_{i=0}^{m} (\|x_i - D\theta_i\|_2 + \lambda \|\theta_i\|_1)$. Traditional approaches rely on heuristic methods such as EM. Several Dictionary Learning algorithms work by iterating the following two steps as in [17, 13, 12]:

1. Solve the Dictionary Learning problem for all data points X. This can be done using your favorite vector selection algorithm, such as basis pursuit from [3].

2. Given X, the optimization problems is now convex in D.

Let $X = [x_1 \dots x_m]$ and $\Theta = [\theta_1 \dots \theta_m]$. Using a maximum likelihood formalism, the Method of Optimal Dictionary (MOD) [4] uses the pseudoinverse to compute D: $D^{(i+1)} = X{\Theta^{(i)}}^T (\Theta^n {\Theta^i}^T)^{-1}$. The MOD can be extended to Maximum A-Posteriori probability setting with different priors to take into account preferences in the recovered dictionary. Similarly, k-SVD uses a two step iterative process, with a Truncated Singular Value Decomposition to update D. This is done by taking every atom in D and applying SVD to X and Θ restricted to only the columns that that have contribution from that atom. When D is restricted to be of the form $D = [B_1, B_2 \dots B_L]$ where B_i's are orthonormal matrices, a more efficient pursuit algorithm is obtained for the sparse coding stage using a block coordinate relaxation.

2 Main Result: Combinatorial Rigidity Characterization for Geometric Dictionary Learning

Now we prove the main result of the paper, i.e, a complete solution to the problem of finding a dictionary D for data X, when the hypergraph $H(S_{X,D})$ of the underlying

subspace arrangement is specified. Additionally we give a (combinatorial) characterization of the (multi)hypergraphs H such that the existence and local uniqueness of a dictionary D is guaranteed for generic X satisfying $H = H(S_{X,D})$.

Problem 5 (Restricted Dictionary Learning)**.** *Let X be a given set of data points in $\mathbb{R}^d$. For an unknown s-regular dictionary $D = [v_1, \ldots, v_n]$ that s-represents X, we are given the hypergraph $H(S_{X,D})$ of the underlying subspace arrangement $S_{X,D}$. Find any (optionally frame) dictionary $\acute{D}$ such that $|\acute{D}| \leq n$, that is consistent with the hypergraph $H(S_{X,D})$.*

This simplification enables us to analyze the problem using machinery from algebraic and combinatorial geometry.

Since the magnitudes of the vectors in X or D are uninteresting, we treat the data and dictionary points in the projective $(d-1)$ space and use the same notation to refer to both original d-dimensional and projective $d-1$ dimensional versions when the meaning is clear from the context.

Problem 6 (Pinned Subspace-Incidence Problem)**.** *Let X be a given set of m points (pins) in $\mathbb{P}^{d-1}(\mathbb{R})$. For every data point $x \in X$, we are also given the hyperedge $supp_D(x)$, i.e, a subset of an unknown set of points $D = \{v_1, \ldots, v_n\}$, such that x_i lies on the subspace spanned by $supp_D(x)$; assume additionally that no s pins lie on any single subspace defined by one of the above hyperedges. Find any such set of points D that satisfies the given subspace incidences.*

In the following, we prove combinatorial conditions that characterize the class of inputs that recover a finite number of solutions D.

2.1 Algebraic Representation

We derive an algebraic system of equations to represent our problem in the tradition of geometric constrain solving [2, 16]. For convenience, we denote a minor of a matrix A using the notation $A[R, C]$, where R and C are index sets of the rows and columns contained in the minor, respectively. In addition, $A[R, \cdot\,]$ represents the minor containing all columns and row set R, and $A[\,\cdot\,, C]$ represents the minor containing all rows and column set C.

Consider a pin x_k on the subspace spanned by points $v_1^k, v_2^k, \ldots, v_s^k$. Using homogeneous coordinates, we can write this incidence constraint by letting all the $s \times s$ minors of the $(d-1) \times s$ matrix

$$E^k = \begin{bmatrix} v_1^k - x_k & v_2^k - x_k & \ldots & v_s^k - x_k \end{bmatrix}$$

be zero, where $v_i^k = (v_{i,1}^k, v_{i,2}^k, \ldots, v_{i,d-1}^k)$ and $x_k = (x_{k,1}, x_{k,2}, \ldots, x_{k,d-1})$. So each incidence can be written as $\binom{d-1}{s}$ equations:

$$E^k[R(l), \cdot\,] = 0, \qquad 1 \leq l \leq \binom{d-1}{s} \tag{1}$$

where $R(l)$ enumerates all the s-subsets of rows of E^k. Note that only $d-s$ of these $\binom{d-1}{s}$ equations are independent, as the subspace spanned by $v_1^k, v_2^k, \ldots, v_s^k$ is a s-dimensional subspace in a d-dimensional space, which only has $s(d-s)$ degrees of freedom.

Given the hypergraph $H = H(S_{X,D})$ of the underlying subspace arrangement, the pinned subspace-incidence problem now reduces to solving a system of $m\binom{d-1}{s}$ equations (or, equivalently, $m(d-s)$ independent equations), each of the form (1). The system of equations sets a multivariate function $(H,X)(D)$ to 0:

$$(H,X)(D) = \begin{cases} \cdots \\ E^k[R(l),\ \cdot\,] \quad = \quad 0 \\ \cdots \end{cases} \tag{2}$$

When viewing X as a fixed parameter, $(H,X)(D)$ is a vector valued function from $R^{n(d-1)}$ to $R^{m\binom{d-1}{s}}$ parameterized by X.

Without any pins, the points in D have in total $n(d-1)$ degrees of freedom. In general, putting r pins on an s-dimensional subspace of d-dimensional space gives an $(s-r)$-dimensional subspace of a $(d-r)$-dimensional space, which has $(s-r)((d-r)-(s-r)) = (s-r)(d-s)$ degrees of freedom left. So every pin potentially removes $(d-s)$ degrees of freedom.

2.2 Linearization as Rigidity Matrix and its Generic Combinatorics

As shown in the previous section, the Pinned Subspace-Incidence problem can be viewed as finding the common solutions of a system of polynomial equations (finding a real algebraic variety). We describe the approach taken by the traditional rigidity theory [1, 6] for characterizing generic properties of these solutions and give some of the definitions.

We use the underlying (multi)hypergraph $H(S_{X,D}) = (\mathfrak{I}(D), \mathfrak{I}(S_{X,D}))$ to define a *pinned subspace-incidence framework* (H,X,D), where $X : \{x_1,\ldots,x_m\} \subseteq \mathbb{R}^{d-1} \to \mathfrak{I}(S_{X,D})$ is an assignment of a given set of pins x_k to edges $X(x_k) = supp_D(x_k) \in \mathfrak{I}(S_{X,D})$, and $D : \mathfrak{I}(D) \to \mathbb{R}^{d-1}$ is an embedding of each vertex j into a point $v_j \in \mathbb{R}^{d-1}$, such that each pin x_k lies on the subspace spanned by $\{v_1^k, v_2^k, \ldots, v_s^k\}$. Note: when the context is clear, we use X to denote both the set of points $\{x_1,\ldots,x_m\}$, as well as the above assignment of these points to edges of H. Two frameworks (H_1,X_1,D_1) and (H_2,X_2,D_2) are *equivalent* if $H_1 = H_2$ and $X_1 = X_2$, i.e. they satisfy the same algebraic equations for the same labeled hypergraph and ordered set of pins. They are *congruent* if they are equivalent and $D_1 = D_2$.

The pinned subspace-incidence system $(H,X)(D)$ is *independent* if none of the algebraic constraints is in the ideal generated by the others. Generally, independence implies the existence of a solution D to the system $(H,X)(D)$, where X is fixed. The system is *rigid* if there exist at most finitely many (real or complex) solutions. The system is *minimally rigid* if it is both rigid and independent. The system is *globally rigid* if there exists at most one solution. Rigidity and global rigidity are often defined (slightly differently) for individual frameworks. A framework (H,X,D) is *rigid*

(i.e. *locally unique*) if there is a neighborhood $N(D)$ of D, such that any framework (H, X, D') equivalent to (H, X, D) with $D' \in N(D)$ is also congruent to (H, X, D). A rigid framework (H, X, D) is *minimally rigid* if it becomes flexible after removing any pin. A framework (H, X, D) is *globally rigid* (i.e. *globally unique*) if any framework equivalent to (H, X, D) is also congruent to (H, X, D).

Pinned subspace-incidence frameworks are generalizations of related types of frameworks, such as in pin-collinear body-pin frameworks [8], direction networks [24], slider-pinning rigidity [19], the molecular conjecture in 2D [15], body-cad constraint system [7], k-frames [22, 23], and affine rigidity [5].

2.2.1 Genericity

Checking independence relative to the ideal generated by the variety is computationally hard and best known algorithms, such as computing Grobner basis, are exponential in time and space [10]. However, the algebraic system can be linearized at *regular* (non-singular) points whereby independence and rigidity of the algebraic pinned subspace-incidence system $(H, X)(D)$ reduces to linear independence and maximal rank at *generic* frameworks.

In algebraic geometry, a property is generic intuitively means that the property holds on the open dense complement of an (real) algebraic variety. Formally,

Definition 2. *A framework (H, X, D) is generic w.r.t. a property Q if and only if there exists a neighborhood $N(D)$ such that for all frameworks (H, X, D') with $D' \in N(D)$, (H, X, D') satisfies Q if and only if (H, X, D) satisfies Q.*

Furthermore we can define generic properties of the hypergraph.

Definition 3. *A property Q of frameworks is generic (i.e, becomes a property of the hypergraph alone) if for all graphs H, either all generic (w.r.t. Q) frameworks (H, X, D) satisfies Q, or all generic (w.r.t. Q) frameworks (H, X, D) do not satisfy Q.*

A framework (H, X, D) is *generic* for property Q if an algebraic variety V_Q specific to Q is avoided by the given framework (H, X, D). Often, for convenience in relating Q to other properties, a more restrictive notion of genericity is used than stipulated by Definition 2 or 3, i.e. another variety $\acute{V}_Q$ is chosen so that $V_Q \subseteq \acute{V}_Q$, as in Lemma 4. Ideally, the variety $\acute{V}_Q$ corresponding to the chosen notion of genericity should be as tight as possible for the property Q (necessary and sufficient for Definition 2 and 3), and should be explicitly defined, or at least easily testable for a given framework.

Once an appropriate notion of genericity is defined, we can treat Q as a property of a hypergraph. The primary activity of the area of combinatorial rigidity is to give purely combinatorial characterizations of such generic properties Q. In the process of drawing such combinatorial characterizations, the notion of genericity may have to be further restricted, i.e. the variety $\acute{V}_Q$ is further expanded by so-called pure conditions that are necessary for the combinatorial characterization to go through (we will see this in Theorem 8).

2.2.2 Linearization

Adapting [1], we now show that rigidity and independence (based on nonlinear polynomials) of pinned pubspace-incidence systems are generically properties of the underlying hypergraph $H(S_{X,D})$, and can furthermore be captured by linear conditions in an generic infinitesimal setting. A *rigidity matrix* of a framework (H, X, D) is a matrix whose kernel is the infinitesimal motions (flexes) of (H, X, D). A framework is *infinitesimally independent* if the rows of the rigidity matrix are independent. A framework is *infinitesimally rigid* if there the space of infinitesimal motion is trivial, i.e. the rigidity matrix has full rank. A framework is *infinitesimally minimally rigid* if it is both infinitesimally independent and rigid.

Note that the rank of a generic matrix M is at least as large as the rank of any specific realization $M(H, X, D)$.

To define a rigidity matrix for a pinned subspace-incidence framework (H, X, D), we take the Jacobian $J_X(D)$ of the algebraic system $(H, X)(D)$. by taking partial derivatives w.r.t. the coordinate of v_i's. In the Jacobian, each pin x_k has $\binom{d-1}{s}$ corresponding rows, and each vertex v_i has $d-1$ corresponding columns. Each equation $E^k[R(l), \cdot\,] = 0$ (1) gives the corresponding row in the Jacobian:

$$\begin{aligned} r^k(l) = [&0, \ldots, 0, 0, V^k_{1,1}(l), V^k_{1,2}(l), \ldots, V^k_{1,d-1}(l), 0, 0, \\ &\cdots\cdots \\ &\ldots, 0, 0, V^k_{s,1}(l), V^k_{s,2}(l), \ldots, V^k_{s,d-1}(l), 0, 0, \ldots, 0] \end{aligned}$$

Each vertex v^k_i has the entries $V^k_{i,1}(l), V^k_{i,2}(l), \ldots, V^k_{i,d-1}(l)$ in its $d-1$ columns, where s of which are generically non-zero. Let $D^k = \{v^k_i, 1 \le i \le s\}$ be the vertices of the hyperedge corresponding to x_k. An entry $V^k_{i,j}(l)$, where $j \in R(l)$, stands for the $(s-1)$-dimensional volume, of the $(s-1)$-simplex formed by the vertices $(D^k \setminus \{v^k_i\})$ together with x_k, projected on the coordinates $R(l) \setminus \{j\}$. For a generic framework, we can choose a coordinate system such that these projected volumes are non-zero. All the other entries, including the terms $V^k_{i,j}(l)$ where $j \notin R(l)$, and the entries corresponding to vertices not on the hyperedge x_k, are zero.

Notice that for every pair of vertices v^k_i and $v^k_{i'}$, the projected volumes on different coordinates all have the same ratio: $\dfrac{V^k_{i,j_2}(l)}{V^k_{i,j_1}(l)} = \dfrac{V^k_{i',j_2}(l)}{V^k_{i',j_1}(l)}$ for all $1 \le j_1, j_2 \le d-1, j_1 \in R(l), j_2 \in R(l)$. So we can divide each row $r^k(l)$ by $\sum_{i=1}^{s} V^k_{i,j^*}(l)$, where j^* is the smallest index in $R(l)$, and simplify $r^k(l)$ to:

$$\begin{aligned} &[0, \ldots, 0, 0, b_1a_1, b_2a_1 \ldots, b_{d-1}a_1 0, 0, \\ &\quad \cdots\cdots, \\ &\quad \ldots, 0, 0, b_1a_{s-1}, b_2a_{s-1}, \ldots, b_{d-1}a_{s-1}, 0, 0, \\ &\quad \ldots, 0, 0, b_1\left(1 - \sum_{i=1}^{s-1} a_i\right), b_2\left(1 - \sum_{i=1}^{s-1} a_i\right), \ldots, b_{d-1}\left(1 - \sum_{i=1}^{s-1} a_i\right), 0, 0, \ldots, 0] \end{aligned} \tag{3}$$

where the values of a_i and b_j are related to l and k, and $b_j = 0$ if $j \notin R(l)$.

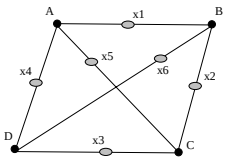

Figure 1: An pinned subspace-incidence framework of 6 pins and 4 vertices, with $d = 4, s = 2$.

Example 1. *Figure 1 shows a pinned subspace-incidence framework with $d = 4, s = 2$. If we denote $\alpha_{1,1} = A_1 - x_{1,1}, \alpha_{1,2} = A_2 - x_{1,2}, \beta_{1,2} = B_2 - x_{1,2}$, etc, the edge AB will have the following three rows in the Jacobian:*

$$\begin{bmatrix} \beta\alpha_{1,2} & -\beta\alpha_{1,1} & 0 & -\alpha_{1,2} & \alpha_{1,1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \beta\alpha_{1,3} & 0 & -\beta\alpha_{1,1} & -\alpha_{1,3} & 0 & \alpha_{1,1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \beta\alpha_{1,3} & -\beta\alpha_{1,2} & 0 & -\alpha_{1,3} & \alpha_{1,2} & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

and the corresponding rows in the simplified Jacobian has the following form

$$\begin{bmatrix} b_{1,1}a_1 & b_{1,2}a_1 & 0 & b_{1,1}(1-a_1) & b_{1,2}(1-a_1) & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ b_{2,1}a_2 & 0 & b_{2,2}a_2 & b_{2,1}(1-a_2) & 0 & b_{2,2}(1-a_2) & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & b_{3,1}a_3 & b_{3,2}a_3 & 0 & b_{3,1}(1-a_3) & b_{3,2}(1-a_3) & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

For a pinned subspace-incidence framework (H, X, D), we define the *symmetric rigidity matrix* M to be the simplified Jacobian matrix obtained above, of size $\binom{d-1}{s}$ by $n(d-1)$, where each row has the form (3). The framework is infinitesimally rigid if and only if M has full column rank.

Notice that in M each hyperedge has $\binom{d-1}{s}$ rows, where any $d - s$ of them are independent, and spans all other rows. If we choose $d - s$ rows per hyperedge in M, the obtained matrix $\hat{M}$ is a rigidity matrix of size $m(d-s)$ by $n(d-1)$.

Remark 1. *There are several correct ways to write the rigidity matrix of a framework, depending on what one considers as the primary indeterminates (points, subspaces, or both), i.e. whether one chooses to work in primal or dual space. We pick points for columns for the simplicity of the row pattern.*

Adapting [1], and defining generic as non-singular, we show that for a generic framework (H, X, D), infinitesimal rigidity is equivalent to generic rigidity.

Lemma 4. *If D and X are regular / non-singular with respect to the system $(H, X)(D)$, then generic infinitesimal rigidity of the framework (H, X, D) is equivalent to generic rigidity.*

Proof Sketch. First we show that if a framework is regular, infinitesimal rigidity implies rigidity. Consider the polynomial system $(H, X)(D)$ of equations. The Implicit Function Theorem states that there exists a function g, such that $D = g(X)$ on some

open interval, if and only if the Jacobian $J_X(D)$ of $(H, X)(D)$ with respect to D has full rank. Therefore, if the framework is infinitesimally rigid, then the solutions to the algebraic system are isolated points (otherwise g could not be explicit). Since the algebraic system contains finitely many components, there are only finitely many such solution and each solution is a 0 dimensional point. This implies that the total number of solutions is finite, which is the definition of rigidity.

To show that generic rigidity implies generic infinitesimal rigidity, we take the contrapositive: if the framework is not infinitesimally rigid, we show that there is a finite flex. If (H, X, D) is not infinitesimally rigid, then the rank r of the Jacobian $J_X(D)$ is less than $2m$. Let E^* be a set of edges in H such that $|E^*| = r$ and the corresponding rows in the Jacobian $J_X(D)$ are all independent. There are r independent rows as well. Let D_{E^*} be the components of D corresponding to those r rows and $D_{E^{*\perp}}$ be the remaining components. The r-by-r submatrix, made of up of the corresponding independent rows and columns, is invertible. Then, by the Implicit Function Theorem, in a neighborhood of D there exists a continuous and differentiable function g such that $D_{E^*} = g(D_{E^{*\perp}})$. This identifies D^*, whose components are D_{E^*} and the level set of g corresponding to D_{E^*}, such that $(H, X)(D^*) = 0$. The level set defines the finite flexing of the framework. Therefore the system is not rigid. □

2.3 Required Hypergraph Properties

This section introduces a pure hypergraph property that will be useful for our characterization.

Definition 5. *A hypergraph $H = (V, E)$ is $(k, 0)$-sparse if for any $V' \subset V$, the induced subgraph $H' = (V', E')$ satisfies $|E'| \leq k|V'|$. A hypergraph H is $(k, 0)$-tight if H is $(k, 0)$-sparse and $|E| = k|V|$.*

This is a special case of the (k, l)-sparsity condition that was formally studied widely in the geometric constraint solving and combinatorial rigidity literature before it was given a name in [9]. A relevant concept from graph matroids is *map-graph*, defined as follows.

Definition 6. *An* orientation *of a hypergraph is given by identifying as the* tail *of each edge one of its endpoints. The* out-degree *of a vertex is the number of edges which identify it as the tail and connect v to $V - v$. A* map-graph *is a hypergraph that admits an orientation such that the out degree of every vertex is exactly one.*

The following lemma follows Tutte-Nash Williams [20, 11] to give a useful characterization of $(k, 0)$-tight graphs in terms of maps.

Lemma 7 ([18])**.** *A hypergraph H is composed of k edge-disjoint map-graphs if and only if H is $(k, 0)$-tight.*

2.4 Main Theorem: Combinatorial Characterization of Dictionary Hypergraphs

We obtain the following combinatorial characterization of existence of a dictionary (sparsity or independence) and local uniqueness or finite (possibly complex) solution

set (rigidity) for a pinned subspace-incidence framework.

Theorem 8 (Main Theorem)**.** *A pinned subspace-incidence framework is generically minimally rigid if and only if the underlying hypergraph* $H(S_{X,D}) = (\mathfrak{I}(D), \mathfrak{I}(S_{X,D}))$ *satisfies* $(d-s)\|\mathfrak{I}(S_{X,D})\| = (d-1)|\mathfrak{I}(D)|$ *(i.e.* $(d-s)|X| = (d-1)|D|$*), and* $(d-s)|E'| \leq (d-1)|V'|$ *for every vertex induced subgraph* $H' = (V', E')$*. The latter condition alone ensures the independence of the framework.*

The graph property from Theorem 8 is not directly a $(k, 0)$-tightness condition, so we modify the underlying hypergraph by duplicating each hyperedge into $(d-s)$ copies.

Definition 9. *Given a underlying hypergraph* $H = (\mathfrak{I}(D), \mathfrak{I}(S_{X,D}))$ *of a Pinned Subspace-Incidence problem, the* expanded multihypergraph $\hat{H} = (V, \hat{E})$ *of* H *is obtained by letting* $V = \mathfrak{I}(D)$*, and replacing each hyperedge in* $\mathfrak{I}(S_{X,D})$ *with* $(d-s)$ *copies in* $\hat{E}$*.*

Theorem 8 can be restated on the expanded multihypergraph:

Theorem 10. *A pinned subspace-incidence framework is generically minimally rigid if and only if the underlying expanded multihypergraph is* $(d-1, 0)$*-tight.*

Example 2. *Figure 1 gives an example of a framework whose expanded multihypergraph is* $(3, 0)$*-tight.*

Since Theorem 10 is equivalent to Theorem 8, we only need to prove Theorem 10 in the following.

The *rigidity matrix* $\hat{M}$ for a pinned subspace-incidence framework is a $(d-s)|\hat{E}|$ by $(d-1)|V|$ matrix according to the expanded multihypergraph $\hat{H} = (V, \hat{E})$, where each hyperedge $x_k \in \mathfrak{I}(S_{X,D})$ has $(d-s)$ rows, one for each copy. The $(d-s)$ rows are arbitrarily picked from the $\binom{d-1}{s}$ rows of x_k in the symmetric rigidity matrix M.

The proof adopts an approach by [23, 22], in proving rigidity of k-frames. The proof outline is as follows:

- We show that for a specific form of the rows of a matrix defined on a map-graph, the determinant is not identically zero (Lemma 11).
- We apply Laplace decomposition to the $(d-1, 0)$-tight hypergraph as a union of $d-1$ maps, to show that the determinant of the rigidity matrix is not identically zero. (Proof of Main Theorem).
- The resulting polynomial is called the *pure condition*: the relationship that the framework has to satisfy in order for the combinatorial characterization to hold.

We first consider the generic rank of particular matrices defined on a single map-graph.

Lemma 11. *A matrix N defined on a map-graph $H = (V, E)$, such that columns are indexed by the vertices and rows by the edges, where the row for hyperedge $x_k \in E$ has non-zero entries only at the s indices corresponding to $v_i^k \in x_k$, with the following pattern:*

$$[0, \ldots, 0, a_1^k, 0, \ldots, a_2^k, 0, \ldots\ldots, 0, a_{s-1}^k, 0, \ldots, 1 - \sum_{i=1}^{s-1} a_i^k, 0, \ldots, 0] \qquad (4)$$

is generically full rank.

Proof. According to the definition of a map-graph, the function $t : E \to V$ assigning a tail vertex to each hyperedge is a one-to-one correspondence. Without loss of generality, assume that for any x_k, the corresponding entry of $t(x_k)$ in N is $1 - \sum_i a_i^k$ (notice that we can arbitrarily switch the variable names $a_1^k, \ldots, a_{s-1}^k, 1 - \sum_i a_i^k$). The determinant of the map N is:

$$\det(N) = \pm \prod_k (1 - \sum_i a_i^k) + \sum_\sigma \text{sgn}(\sigma) \prod_{i=1}^{n} N[i, \sigma_i] \qquad (5)$$

where σ enumerates all other permutations of $|N|$, excluding that of the first term $\pm \prod_k (1 - \sum_i a_i^k)$.

Notice that each term $\prod_{i=1}^n N[i, \sigma_i]$ has at least one a_i^k as a factor. If we use the specialization with $a_i^k = 0$ for all i and k, the summation over σ will all be zero, and $\det(N)$ will be $\pm \prod_k (1 - \sum_i a_i^k) = \pm 1$. So generically, N must have full rank. □

Now we are ready to prove the main theorem.

Proof of Main Theorem. First we show the only if direction. For a generically minimally rigid pinned subspace-incidence framework, the determinant of $\hat{M}$ is not identically zero. Since the number of columns is $n(d-1)$, it is trivial that $n(d-1)$ copied edges in $\hat{M}$, namely $n\frac{d-1}{d-s}$ pins, are necessary. It is also trivial to see that $(d-1, 0)$-tightness is necessary, since any subgraph $H' = (V', E')$ of $\hat{H}$ with $|E'| > (d-1)|V'|$ is overdetermined and generically has no solution.

Next we show the if direction, that $n(d-1)$ copied edges arranged generically in a $(d-1, 0)$-tight pattern imply infinitesimal rigidity.

We first group the columns according to the coordinates. In other words, we have $d-1$ groups C_j, where all columns for the first coordinate belong to C_1, all columns for the second coordinate belong to C_2, etc. This can be done by applying a Laplace expansion to rewrite the determinant of the rigidity matrix $\hat{M}$ as a sum of products of determinants (brackets) representing each of the coordinates taken separately:

$$\det(\hat{M}) = \sum_\sigma \left(\pm \prod_j \det \hat{M}[R_j^\sigma, C_j] \right)$$

where the sum is taken over all partitions σ of the rows into $d-1$ subsets $R_1^\sigma, R_2^\sigma, \ldots, R_j^\sigma, \ldots, R_{d-1}^\sigma$, each of size $|V|$. Observe that for each $\hat{M}[R_j^\sigma, C_j]$,

$$\det(\hat{M}[R_j^\sigma, C_j]) = (b_1^{\sigma j} \ldots b_n^{\sigma j}) \det(M'[R_j^\sigma, C_j])$$

for some coefficients $(b_1^{\sigma j} \dots b_n^{\sigma j})$, and each row of $\det(M'[R_j^\sigma, C_j])$ is either all zero, or of pattern (4). By Lemma 7, the expanded multihypergraph $\hat{H}$ can be decomposed into $(d-1)$ edge-disjoint maps. Each such decomposition has some corresponding row partitions σ, where each column group C_j corresponds to a map N_j, and R_j^σ contains rows corresponding to the edges in that map. Observe that $\hat{M}[R_j^\sigma, C_j]$ contains an all-zero row r, if and only if the row r has the jth coordinate entry being zero in $\hat{M}$. Recall for each hyperedge x_k, we are free to pick any $d-s$ rows to include in $\hat{M}$ from the $\binom{d-1}{s}$ rows in the symmetric rigidity matrix M. We claim that

Claim 1. *Given a map decomposition, we can always pick the rows of the rigidity matrix $\hat{M}$, such that there is a corresponding row partition σ^*, where none of the minors $\hat{M}[R_j^{\sigma^*}, C_j]$ contains an all-zero row.*

Given a map decomposition, for any map N_j, there are $\binom{d-2}{s-1}$ among these $\binom{d-1}{s}$ rows with the jth coordinate being non-zero. Also, it is not hard to show that for all $2 \le s \le d-1$, $\binom{d-2}{s-1} \ge d-s$. So for any N_j containing k_j copies of a particular hyperedge, since all the other maps can pick at most $(d-s)-k_j$ rows from its $\binom{d-2}{s-1}$ choices, it still has $\binom{d-2}{s-1} - ((d-s)-k_j) \ge k_j$ choices. Therefore, given a map decomposition, we can always pick the rows in the rigidity matrix $\hat{M}$, such that there is a partition of each hyperedge's rows, where each map N_j get its required rows with non-zeros at coordinate j. This concludes the proof of the claim.

So by Lemma 11, the determinate of each such minor $\hat{M}[R_j^{\sigma^*}, C_j]$ is generically non-zero. We conclude that

$$\det(\hat{M}) = \sum_\sigma \left(\pm \prod_j \left((b_1^{\sigma j} \dots b_n^{\sigma j}) \det\left(M'[R_j^\sigma, C_j]\right) \right) \right)$$

Observe that each term of the sum has a unique multi-linear coefficient $(b_1^{\sigma j} \dots b_n^{\sigma j})$ that generically do not cancel with any of the others since $\det(M'[R_j^\sigma, C_j])$ are independent of the b's. This implies that $\hat{M}$ is generically full rank, thus completes the proof. Moreover, substituting the values of $\det(M'[R_j^\sigma, C_j])$ from Lemma 11 gives the pure condition for genericity. □

Example 3. *Consider the pinned subspace-incidence framework in Example 1. The expanded multihypergraph satisfies $(3,0)$-tightness condition. The rigidity matrix $\hat{M}$ has the following form:*

$$\begin{bmatrix}
b_{1,1}a_1 & b_{1,2}a_1 & 0 & b_{1,1}(1-a_1) & b_{1,2}(1-a_1) & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
b_{2,1}a_2 & 0 & b_{2,2}a_2 & b_{2,1}(1-a_2) & 0 & b_{2,2}(1-a_2) & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & b_{3,1}a_3 & b_{3,2}a_3 & 0 & b_{3,1}(1-a_3) & b_{3,2}(1-a_3) & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & b_{4,1}a_4 & 0 & b_{4,2}a_4 & b_{4,1}(1-a_4) & 0 & b_{4,2}(1-a_4) & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & b_{5,1}a_5 & b_{5,2}a_5 & 0 & b_{5,1}(1-a_5) & b_{5,2}(1-a_5) & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & b_{6,1}a_6 & 0 & b_{6,2}a_6 & b_{6,1}(1-a_6) & 0 & b_{6,2}(1-a_6) \\
b_{7,1}a_7 & b_{7,2}a_7 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & b_{7,1}(1-a_7) & b_{7,2}(1-a_7) & 0 \\
b_{8,1}a_8 & 0 & b_{8,2}a_8 & 0 & 0 & 0 & 0 & 0 & 0 & b_{8,1}(1-a_8) & 0 & b_{8,2}(1-a_8) \\
b_{9,1}a_9 & b_{9,2}a_9 & 0 & 0 & 0 & 0 & b_{9,1}(1-a_9) & b_{9,2}(1-a_9) & 0 & 0 & 0 & 0 \\
b_{10,1}a_{10} & 0 & b_{10,2}a_{10} & 0 & 0 & 0 & b_{10,1}(1-a_{10}) & 0 & b_{10,2}(1-a_{10}) & 0 & 0 & 0 \\
0 & 0 & 0 & b_{11,1}a_{11} & b_{11,2}a_{11} & 0 & 0 & 0 & 0 & b_{11,1}(1-a_{11}) & b_{11,2}(1-a_{11}) & 0 \\
0 & 0 & 0 & b_{12,1}a_{12} & 0 & b_{12,2}a_{12} & 0 & 0 & 0 & b_{12,1}(1-a_{12}) & 0 & b_{12,2}(1-a_{12})
\end{bmatrix}$$

After grouping the coordinates, it becomes

$$\left[\begin{array}{cccc|cccc|cccc}
b_{1,1}a_1 & b_{1,1}(1-a_1) & 0 & 0 & b_{1,2}a_1 & b_{1,2}(1-a_1) & 0 & 0 & 0 & 0 & 0 & 0 \\
b_{2,1}a_2 & b_{2,1}(1-a_2) & 0 & 0 & 0 & 0 & 0 & 0 & b_{2,2}a_2 & b_{2,2}(1-a_2) & 0 & 0 \\
0 & b_{3,1}a_3 & b_{3,1}(1-a_3) & 0 & 0 & b_{3,2}a_3 & b_{3,2}(1-a_3) & 0 & 0 & 0 & 0 & 0 \\
0 & b_{4,1}a_4 & b_{4,1}(1-a_4) & 0 & 0 & 0 & 0 & 0 & 0 & b_{4,2}a_4 & b_{4,2}(1-a_4) & 0 \\
0 & 0 & b_{5,1}a_5 & b_{5,1}(1-a_5) & 0 & 0 & b_{5,2}a_5 & b_{5,2}(1-a_5) & 0 & 0 & 0 & 0 \\
0 & 0 & b_{6,1}a_6 & b_{6,1}(1-a_6) & 0 & 0 & 0 & 0 & 0 & 0 & b_{6,2}a_6 & b_{6,2}(1-a_6) \\
b_{7,1}a_7 & 0 & 0 & b_{7,1}(1-a_7) & b_{7,2}a_7 & 0 & 0 & b_{7,2}(1-a_7) & 0 & 0 & 0 & 0 \\
b_{8,1}a_8 & 0 & 0 & b_{8,1}(1-a_8) & 0 & 0 & 0 & 0 & b_{8,2}a_8 & 0 & 0 & b_{8,2}(1-a_8) \\
b_{9,1}a_9 & 0 & b_{9,1}(1-a_9) & 0 & b_{9,2}a_9 & 0 & b_{9,2}(1-a_9) & 0 & 0 & 0 & 0 & 0 \\
b_{10,1}a_{10} & 0 & b_{10,1}(1-a_{10}) & 0 & 0 & 0 & 0 & 0 & b_{10,2}a_{10} & 0 & b_{10,2}(1-a_{10}) & 0 \\
0 & b_{11,1}a_{11} & 0 & b_{11,1}(1-a_{11}) & 0 & b_{11,2}a_{11} & 0 & b_{11,2}(1-a_{11}) & 0 & 0 & 0 & 0 \\
0 & b_{12,1}a_{12} & 0 & b_{12,1}(1-a_{12}) & 0 & 0 & 0 & 0 & 0 & b_{12,2}a_{12} & 0 & b_{12,2}(1-a_{12})
\end{array}\right]$$

where the red rows inside each column group corresponding to a map decomposition of the expanded multihypergraph.

Theorem 8 gives a pure condition that characterizes the badly behaved cases (i.e. the conditions of non-genericity that breaks the combinatorial characterization of the infinitesimal rigidity). The pure condition is a function of the a's and b's which can be calculated from the particular realization (framework) using Lemma 11 and the main theorem. Whether it is possible to efficiently test for genericity from the problem's input (the hypergraph and x_k's) is an open problem.

One particular situation avoided by the pure condition is that there cannot be more than $s-1$ hyperedges containing the same set of vertices, namely, more than $s-1$ pins on the same subspace spanned by the dictionary vectors. This is important, otherwise, s pins completely determine an s-subspace, whereby the vertices of the corresponding hyperedge have their degrees of freedom restricted and simple counterexamples to the characterization of the main theorem can be constructed.

Example 4. *Consider the framework in Figure 2. There are $s = 2$ pins on each subspace. The expanded multihypergraph of the framework is $(2,0)$-tight. However, the framework is obviously not rigid.*

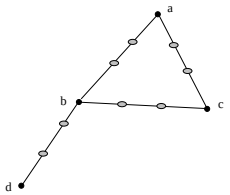

Figure 2: An pinned subspace-incidence framework of 8 pins and 4 vertices, with $d = 3, s = 2$.

Theorem 8 requires the following genericities:

- The pure condition, which is a function of a given framework.
- Generic infinitesimal rigidity, which is the generic rank of the matrix.

The relationship between the two notions of genericities is an open question. Whether one implies the other is an area of future development. However, each of the above conditions applies to an open and dense set. Therefore the notion of genericity for the entire theorem that satisfies all of the above conditions is also open and dense.

2.5 Consequences and Dictionary Construction Algorithm

We related the Restricted Dictionary Learning problem to the general Geometric Dictionary Learning problem. The following is a useful corollary to the main theorem.

Corollary 12. *Given a set of m points $X = \{x_1, .., x_m\}$ in $\mathbb{R}^d$, generically there is a dictionary D of size n such that $D\theta_i = x_i$ and $\|\theta_i\|_0 \leq s$ only if $(d-s)m \leq (d-1)n$. Conversely, if $(d-s)m = (d-1)n$ and the supports of x_i (the nonzero entries of the θ_i) are known to form a $(d-1, 0)$-tight hypergraph H, then generically, there is at least one and at most finitely many such dictionaries.*

Proof. We first prove one direction, that there is generically no dictionary of size $|D| = n$ if $(d-s)m > (d-1)n$. For any hypothetical s-subspace arrangement $S_{X,D}$, the expanded multihypergraph $\hat{H}(S_{X,D})$ - with the given bound for $|D|$ - cannot be $(d-1, 0)$-sparse. Hence generically, under the pure conditions of Theorem 8, the rigidity matrix - of the s-subspace framework $H(S_{X,D})$ - with indeterminates representing the coordinate positions of the points in D - has dependent rows. In which case, the original algebraic system $(H, X)(D)$(whose Jacobian is the rigidity matrix) will not have a (complex or real) solution for D, with X plugged in.

The converse is implied from our theorem since we are guaranteed both generic independence (the existence of a solution) and generic rigidity (at most finitely many solutions). □

Next we quantify the term "generically" in Corollary 12, yielding Corollaries 14 and 13 below.

Corollary 13 (Straightforward Algorithm Constructing Dictionary for Highly General Data Points)**.** *Given a set of m points $X = [x_1 \dots x_m]$ picked uniformly at random from the sphere S^{d-1}, we have a straightforward algorithm to construct a dictionary $D = [v_1 \dots v_n]$ that s-represents X, where $n = \left(\frac{d-s}{d-1}\right) m$.*

Algorithm for constructing the underlying hypergraph $H(S_{X,D})$ for a hypothetical s-subspace arrangement $S_{X,D}$:

The algorithm works in two stages to construct a expanded multihypergraph $\hat{H}(S_{X,D})$:

1. We construct a minimal minimally rigid hypergraph $H_0 = (V_0, E_0)$, using the *pebble game algorithm* introduced below. Here $|V_0| = k(d-s)$, $|E_0| = k(d-1)$, where k is the smallest integer ensuring that no more than $s-1$ edges contain the same set of vertices, i.e. $\binom{k(d-s)}{s}(s-1) \geq k(d-1)$. We assume that $|E_0| << m$.

2. Each construction step adds a basic structure B, which appends $d-s$ vertices and $d-1$ edges to the constructed hypergraph. This B can be obtained using the pebble game algorithm.

The *pebble game algorithm* by [18] works on a fixed finite set V of vertices and constructs a (k,l)-sparse hypergraph. Conversely, any (k,l)-sparse hypergraph on vertex set V can be constructed by this algorithm. This algorithm initializes by putting k pebbles on each vertex in V. There are two types of moves:

- *Add-edge*: adds an hyperedge e (e must contain at least $l+1$ pebbles), removes a pebble from a vertex v in e, and assign v as the tail of e;

- *Pebble-shift*: for a vertex $v_1 \in e$ which contains at least one pebble, let v_2 be the tail of e, move one pebble from v_1 to v_2, and change the tail of e to v_1.

At the end of the algorithm, if there are exactly l pebbles in the hypergraph, then the hypergraph is (k,l)-tight.

Our algorithm runs a slightly modified pebble game algorithm to find a $(d-1,0)$-tight expanded multihypergraph. We require that each add-edge move adding $(d-s)$ copies of a hyperedge e, so a total of $d-s$ pebbles are removed from vertices in e. Additionally, the multiplicity of a hyperedge (not counting the copies) cannot exceed $(s-1)$. To construct the basic structure B of Stage 2, the algorithm initializes by putting $d-1$ pebbles on each vertex in B, and an add-edge move can only add a hyperedge that contains at least one newly added vertex.

Given X, d and s, the underlying hypergraph will have $m = |X|$ edges and $n = m\left(\frac{d-s}{d-1}\right)$ vertices, so Stage 2 will be repeated $(m - |E_0|)/(d-1)$ times. The time complexity for constructing the underlying hypergraph is the time for constructing the framework on H_0, plus the time for constructing the hypergraph structure of B, plus $(m - |E_0|)/(d-1)$ timed by a constant.

Algorithm for constructing the s-subspace arrangement $S_{X,D}$ and the dictionary D:

The construction of the s-subspace arrangement $S_{X,D}$ naturally follows from the construction of the underlying hypergraph. For the initial hypergraph H_0, we get a pinned subspace-incidence system $(H_0, X_0)(D_0)$ by arbitrarily choose $|X_0| = |E_0|$ pins from X. Similarly, for each iteration of Step 2, we form a pinned subspace-incidence system $(B, X_i)(D_i)$ by arbitrarily choosing $|X_i| = d - 1$ pins from X. Notice that both H_0 and B - when positions of its neighbor vertices in the constructed hypergraph are fixed - are minimally rigid.

Given X_i, we know that the rigidity matrix - of the s-subspace framework $B(S_{X_i,D_i})$ - with indeterminates representing the coordinate positions of the points in D_i - generically has full rank (rows are maximally independent), under the pure conditions of Theorem 8; in which case, the original algebraic subsystem $(B, X_i)(D_i)$ (whose Jacobian is the rigidity matrix), with X_i plugged in, is guaranteed to have a (possibly complex) solution and only finitely many solutions for D_i. Since the pure conditions fail only on a measure-zero subset of the space of pin-sets X_i, where each pin is in S^{d-1}, it follows that if the pins in X_i are picked uniformly at random from S^{d-1} we know such a solution exists for D_i (and S_{X_i,D_i}) and can be found by solving the algebraic system $B(S_{X_i,D_i})$, which is of constant size. Similarly, there exists a solution D_0 by solving the subsystem $(H_0, X_0)(D_0)$. We obtain D by taking the union of all D_i's.

Given X, d and s, the time complexity for constructing $S_{X,D}$ is the time for solving the algebraic system $(H_0, X_0)(D_0)$plus $(m-|E_0|)/(d-1)$ timed by the constant time for solving the system $(B, X_i)(D_i)$.

Quantifying "genericity" in Corollary 12 also gives a lower bound on $|D|$.

Corollary 14. *Given a set of m points $X = \{x_1, .., x_m\}$ picked uniformly at random from the sphere S^{d-1}, the corresponding dictionary will have size $\left(\frac{d-s}{d-1}\right) m$ with probability 1. In other words, $|D| = \Omega(X)$ if s and d are constants.*

Proof. To quantify the term "generically" in Corollary 12, we note that the pure-conditions fail only on a measure-zero subset of the space of frameworks $S_{X,D}$. Since the number of possible multi-hypergraphs representing the s-subspace arrangements is finite for a given set of pins, it follows that except for a measure-zero subset of the space of pin-sets X, there is no (real or complex) solution to the algebraic system $(H, X)(D)$. Thus when X is picked uniformly at random from the sphere S^{d-1}, if $|D|$ is less than $\big((d-1)/(d-s)\big)|X|$, with probability 1, there is no solution $S_{X,D}$. □

3 Conclusion

In this paper, we approached Dictionary Learning from a geometric point of view.

We established formal relationships with several related problem of independent interest, which led to a new class of algorithms for Geometric Dictionary Learning.

We investigated Geometric Dictionary Learning theoretically using machinery from algebraic and combinatorial geometry. Specifically, a rigidity-type theorem (our main result) is obtained which completely characterizes the subspace arrangements that are

guaranteed to recover a finite number of dictionaries, using a purely combinatorial property on the supports.

As corollaries of the main result, we gave lower bound for the size of dictionary when the data points are picked uniformly at random, and provided an algorithm for Geometric Dictionary Learning for such general data points.

Additionally, we discussed several open questions leading to different directions for future work.

References

[1] Leonard Asimow and Ben Roth. The rigidity of graphs. *Transactions of the American Mathematical Society*, 245:279–289, 1978.

[2] Beat Brüderlin and Dieter Roller. *Geometric constraint solving and applications*. Springer, 1998.

[3] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM journal on scientific computing*, 20(1):33–61, 1998.

[4] Kjersti Engan, Sven Ole Aase, and J Hakon Husoy. Method of optimal directions for frame design. In *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, volume 5, pages 2443–2446. IEEE, 1999.

[5] Steven J Gortler, Craig Gotsman, Ligang Liu, and Dylan P Thurston. On affine rigidity. *Journal of Computational Geometry*, 4(1):160–181, 2013.

[6] Jack E Graver, Brigitte Servatius, and Herman Servatius. *Combinatorial rigidity*, volume 2. AMS Bookstore, 1993.

[7] Kirk Haller, Audrey Lee-St John, Meera Sitharam, Ileana Streinu, and Neil White. Body-and-cad geometric constraint systems. *Computational Geometry*, 45(8):385–405, 2012.

[8] Bill Jackson and Tibor Jordán. Pin-collinear body-and-pin frameworks and the molecular conjecture. *Discrete & Computational Geometry*, 40(2):258–278, 2008.

[9] Audrey Lee, Ileana Streinu, and Louis Theran. Graded sparse graphs and matroids. *J. UCS*, 13(11):1671–1679, 2007.

[10] Johannes Mittmann. Gröbner bases: Computational algebraic geometry and its complexity. 2007.

[11] C St JA Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 1(1):445–450, 1961.

[12] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research*, 37(23):3311–3325, 1997.

[13] Ignacio Ramirez, Pablo Sprechmann, and Guillermo Sapiro. Classification and clustering via dictionary learning with structured incoherence and shared features. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3501–3508. IEEE, 2010.

[14] Geo rey Davis. *Adaptive nonlinear approximations*. PhD thesis, Citeseer, 1994.

[15] Brigitte Servatius. Molecular conjecture in 2D, 2006. 16th Fall Workshop on Computational and Combinatorial Geometry.

[16] Meera Sitharam. Combinatorial approaches to geometric constraint solving: Problems, progress, and directions. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 67:117, 2005.

[17] Pablo Sprechmann and Guillermo Sapiro. Dictionary learning and sparse coding for unsupervised clustering. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 2042–2045. IEEE, 2010.

[18] Ileana Streinu and Louis Theran. Sparse hypergraphs and pebble game algorithms. *European Journal of Combinatorics*, 30(8):1944–1964, 2009.

[19] Ileana Streinu and Louis Theran. Slider-pinning rigidity: a maxwell–laman-type theorem. *Discrete & Computational Geometry*, 44(4):812–837, 2010.

[20] William T Tutte. On the problem of decomposing a graph into n connected factors. *Journal of the London Mathematical Society*, 1(1):221–230, 1961.

[21] René Vidal. Subspace clustering. *Signal Processing Magazine, IEEE*, 28(2):52–68, 2011.

[22] Neil White and Walter Whiteley. The algebraic geometry of motions of bar-and-body frameworks. *SIAM Journal on Algebraic Discrete Methods*, 8(1):1–32, 1987.

[23] Neil L White and Walter Whiteley. The algebraic geometry of stresses in frameworks. *SIAM Journal on Algebraic Discrete Methods*, 4(4):481–511, 1983.

[24] Walter Whiteley. Some matroids from discrete applied geometry. *Contemporary Mathematics*, 197:171–312, 1996.